

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 1 (Currently Amended) A method for constructing and caching a chain of file identifiers in a computing system environment, the chain of file identifiers representing a full path to a file system resource the method comprising the steps of:

processing a file system resource's defined name (DN) into a file identifier (FID) and defined name database;

retrieving from the defined name database a first file identifier for a first file system resource, the retrieved file identifier corresponding to a defined name of the first file system resource;

adding the retrieved first file identifier to a chain of file identifiers, the added file identifier being the first file identifier in the chain;

retrieving a next file identifier for a next file system resource in a full path of the first file system resource;

adding the retrieved next file identifier to the chain; and

repeating said step of retrieving a next file identifier for a next file system resource and said step of adding the next retrieved file identifier to the chain until a file identifier for each file system resource in a full path of the first file system resource is added in the chain;

retrieving a file identifier corresponding to the file system resource, the file system resource being target of an access attempt and having a file identifier chain representing the full path directory of the target system resource;

searching for a security classification category and defined name for the target resource file identifier;

updating the security classification system, when said search finds a security classification category for the target resource file identifier;

determining whether operations for the target file system resource could affect the file system name space; and

terminating constructing and caching of a chain of file identifiers when there is a determination that the operations for the target file system resource do not affect the file system name space.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 2 (Canceled)

Claim 3 (Previously presented) The method as described in claim 2 further comprising the step of flushing a file identifier chain cache when there is a determination that desired operations on the first file system resource could affect the file system name space.

Claim 4 (Previously presented) The method as described in claim 2 further comprising before said file identifier (FID) retrieval step the step of processing a system resources' defined name (DN) and security classification category into a mapping database that holds a FID to DN mapping.

Claim 5 (Previously presented) The method as described in claim 4 wherein said processing system resources step further comprises:

- providing the defined name and security classification category as inputs;
- obtaining a file identifier (FID) for the defined name; and
- adding the FID to DN mapping containing the security classification category to the mapping database.

Claim 6 (Previously presented) The method as described in claim 4 wherein said searching step comprises:

- searching the FID to DN mapping database for the security classification category for the FID of the target resource; and
- returning the security classification category and defined name for a target FID, when a security classification category for the target FID was found during said search.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 7 (Previously presented) The method as described in claim 4 wherein said searching step comprises:

searching the FID to DN mapping database for the security classification category for the FID of the target resource;

retrieving a FID from a constructed FID chain, when the search does not find a security classification category for the FID of the target resource;

searching the FID to DN mapping database for the security classification category for the retrieval FID of the constructed FID chain; and

returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

Claim 8 (Previously presented) The method as described in claim 7 further comprising the steps of:

determining whether there are more entries in the constructed FID chain, when the search does not find a security classification category for the FID used in the search;

retrieving a next FID in the constructed FID chain; and

searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

Claim 9 (Original) The method as described in claim 8 further comprising the step of terminating the method when no security classification category is found for any FID in the FID chain.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 10 (Previously presented) The method as described in claim 3 wherein said flushing step comprises:

- retrieving a path name for the target resource, said path name being to a directory for the target resource;
- obtaining a vnode for the directory;
- generating a FID for the directory using the vnode;
- searching for FID chain matching directory FID; and
- removing FID chain from cache, when matching FID chain is found.

Claim 11 (Previously presented) The method as described in claim 10 further comprising before said searching step the step of sorting a FID chains in a FID chain cache into hash list.

Claim 12 (Previously presented) The method as described in claim 11 wherein said sorting step comprises:

- retrieving a first FID chain in the FID chain list;
- comparing each FID in the first FID chain to a directory FID;
- determining whether there are more FID chains in the list, when said FID chain did not match the directory FID;
- retrieving a next FID chain in the FID, and
- returning to said comparing step using newly retrieved FID chain.

Claim 13 (Previously presented) The method as described in claim 11 wherein said searching step comprises:

- retrieving a first FID chain in the FID chain list;
- comparing each FID in the first FID chain to the directory FID;
- determining whether there are more FID chains in the list, when said FID chain did not match the directory FID; and
- terminating method when no FID chain is found.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 14 (Currently Amended) A computer program product in a computer readable medium for use in constructing and caching a chain of file identifiers in a computing system environment, the chain of file identifiers representing a full path to a file system resource comprising:

instructions for processing a file system resource's defined name (DN) into a file identifier (FID) and defined name database;

instructions for retrieving from the defined name database a first file identifier for a first file system resource, the retrieved file identifier corresponding to a defined name of the first file system resource;

instructions for adding the retrieved first file identifier to a chain of file identifiers, the added first file identifier being the first file identifier in the chain;

instructions for retrieving a next file identifier for a next file system resource in a full path of the first file system resource;

instructions for adding the retrieved next file identifier to the chain; and

instructions for repeating said instructions for retrieving a next file identifier for a next file system resource and said instructions for adding the next retrieved file identifier to the chain, until a file identifier for each system resource in a full path of the first file system resource is added in the chain.

instructions for retrieving a file identifier corresponding to the file system resource, the file system resource being a target of an access attempt and having a file identifier chain representing the full path directory of the target system resource;

instructions for searching for a security classification category and defined name for the target resource file identifier;

instructions for updating the security classification system, when said search finds a security classification category for the target resource file identifier;

instructions for determining whether operations for the target file system resource could affect the file system name space; and

instructions for terminating constructing and caching of a chain of file identifiers when there is a determination that the operations for the target file system resource do not affect the file system name space; and

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

instructions for flushing the a file identifier chain cache when there is a determination that desired operations on the target file system resource could affect the file system name space.

Claim 15 (Canceled)

Claim 16 (Currently Amended) The computer program product as described in claim 14 15 wherein said flushing instructions comprise:

- instructions for retrieving a path name for the target resource, said path name being to a directory for the target resource;
- instructions for obtaining a vnode for the directory;
- instructions for generating a FID for the directory using the vnode;
- instructions for searching for FID chain matching directory FID; and
- instructions for removing FID chain from cache, when matching FID chain is found.

Claim 17 (Currently Amended) The computer program product as described in claim 14 15 wherein said searching instruction further comprises:

- instructions for searching the FID to DN mapping database for the security classification category for the FID of the target resource;
- instructions for retrieving a FID from a constructed FID chain, when the search does not find a security classification category for the FID of the target resource;
- instructions for searching the FID to DN mapping database for the security classification category for the retrieved FID of the constructed FID chain; and
- instructions for returning the security classification category and defined name for the target FID, when a security classification category for the target FID was found during said search.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 18 (Previously presented) The computer program product as described in claim 17 further comprising the steps of:

instructions for determining whether there are more entries in the FID chain, when the search does not find a security classification category for the FID used in the search;

instructions for retrieving a next FID in the constructed FID chain; and

instructions for searching the FID to DN mapping database for the security classification category for the currently retrieved FID of the FID chain.

Claim 19 (Previously presented) The computer program product as described in claim 18 further comprising before said searching, instructions for sorting FID chains in a FID chain cache into hash list.

Claim 20 (Currently Amended) The computer program product as described in claim 14 ~~15~~ wherein said flushing instructions comprise:

instructions for retrieving a path name for the target resource, said path name being to a directory for the target resource;

instructions for obtaining a vnode for the directory;

instructions for generating a FID for the directory using the vnode;

instructions for searching for FID chain matching directory FID; and

instructions for removing FID chain from cache, when matching FID chain is found.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 21 (Currently Amended) The method as described in claim 1 2 wherein said file identifier retrieval step comprises:

- retrieving the path name of the file resource, the file resource being a target of the access attempt;
- obtaining a FID for target resource with said path name;
- determining whether obtained FID is in a FID chain; and
- returning the target FID and FID chain, when the target resource FID was found in a FID Chain Cache.

Claim 22 (Canceled)

Claim 23 (Previously presented) The method as described in claim 21 wherein said file identifier retrieval step comprises:

- retrieving the path name of the file resource, the file resource being target of the access attempt;
- obtaining a FID for target resource with the path name;
- determining whether the obtained FID is in a FID chain; and
- constructing a FID chain for a parent directory for the obtained FID, when no FID chain for the obtained FID is found.

Claim 24 (Previously presented) The method as described in claim 23 wherein said FID chain construction comprises:

- setting a temporary vnode to equal the vnode for the parent of the target resource;
- determining whether the temporary vnode is the root directory; and
- inserting a FID chain into a FID chain cache with the first FID in the chain serving as the entry search key, when temporary vnode is the root directory.

Appl. No. 09/833,346
Amdt. dated August 1, 2006
Reply to Office action of July 25, 2006

Claim 25 (Previously presented) The method as described in claim 23 wherein said FID chain construction step further comprises:

- setting a temporary vnode to equal a vnode for a parent of a target resource;
- determining whether the temporary vnode is a root directory;
- retrieving a vnode for the next parent in a directory path and determining whether that parent is the root directory; and
- repeating said retrieving step until parent is the root of the directory.

Claim 26 (Original) The method as described in claim 25 further comprising the step of inserting a completed FID chain into the FID chain cache when the parent is the root directory.

Claim 27 (Canceled)